

EFFICIENT MACROBLOCK HEADER CODING FOR VIDEO COMPRESSION

TECHNICAL FIELD

5 The invention relates to video coding, and specifically, to an improved method for coding block parameters used in frame-based and object-based video coding formats.

BACKGROUND

10 Full-motion video displays based upon analog video signals have long been available in the form of television. With recent advances in computer processing capabilities and affordability, full-motion video displays based upon digital video signals are becoming more widely available. Digital video systems can provide significant improvements over conventional analog video systems in creating, modifying, transmitting, storing, and playing full-motion video sequences.

15 Digital video displays include large numbers of image frames that are played or rendered successively at frequencies of between 30 and 75 Hz. Each image frame is a still image formed from an array of pixels based on the display resolution of a particular system. As examples, VHS-based systems have display resolutions of 320x480 pixels, NTSC-based systems have display resolutions of 720x486 pixels, and high-definition
20 television (HDTV) systems under development have display resolutions of 1360x1024 pixels.

 The amounts of raw digital information included in video sequences are massive. Storage and transmission of these amounts of video information is infeasible with conventional personal computer equipment. Consider, for example, a digitized form of a
25 relatively low resolution VHS image format having a 320x480 pixel resolution. A full-length motion picture of two hours in duration at this resolution corresponds to 100 gigabytes of digital video information. By comparison, conventional compact optical disks have capacities of about 0.6 gigabytes, magnetic hard disks have capacities of 1-2 gigabytes, and compact optical disks under development have capacities of up to 8
30 gigabytes.

 To address the limitations in storing or transmitting such massive amounts of digital video information, various video compression standards or processes have been

established, including MPEG-1, MPEG-2, and H.26X. These video compression techniques utilize similarities between successive image frames, referred to as temporal or interframe correlation, to provide interframe compression in which motion data and error signals are used to encode changes between frames.

5 In addition, the conventional video compression techniques utilize similarities within image frames, referred to as spatial or intraframe correlation, to provide intraframe compression in which the image samples within an image frame are compressed. Intraframe compression is based upon conventional processes for compressing still images, such as discrete cosine transform (DCT) encoding. This type of coding is
10 sometimes referred to as “texture” or “transform” coding. A “texture” generally refers to a two-dimensional array of image sample values, such as an array of chrominance and luminance values or an array of alpha (opacity) values. The term “transform” in this context refers to how the image samples are transformed into spatial frequency components during the coding process. This use of the term “transform” should be
15 distinguished from a geometric transform used to estimate scene changes in some interframe compression methods.

Interframe compression typically utilizes motion estimation and compensation to encode scene changes between frames. Motion estimation is a process for estimating the motion of image samples (e.g., pixels) between frames. Using motion estimation, the
20 encoder attempts to match blocks of pixels in one frame with corresponding pixels in another frame. After the most similar block is found in a given search area, the change in position of the pixel locations of the corresponding pixels is approximated and represented as motion data, such as a motion vector. Motion compensation is a process for determining a predicted image and computing the error between the predicted image
25 and the original image. Using motion compensation, the encoder applies the motion data to an image and computes a predicted image. The difference between the predicted image and the input image is called the error signal. Since the error signal is just an array of values representing the difference between image sample values, it can be compressed using the same texture coding method as used for intraframe coding of image samples.

30 Although differing in specific implementations, the MPEG-1, MPEG-2, and H.26X video compression standards are similar in a number of respects. The following

description of the MPEG-2 video compression standard is generally applicable to the others.

MPEG-2 provides interframe compression and intraframe compression based upon square blocks or arrays of pixels in video images. A video image is divided into image sample blocks called macroblocks having dimensions of 16 x 16 pixels. In MPEG-2, a macroblock comprises four luminance blocks (each block is 8 x 8 samples of luminance (Y)) and two chrominance blocks (one 8 x 8 sample block each for Cb and Cr).

In MPEG-2, interframe coding is performed on macroblocks. An MPEG-2 encoder performs motion estimation and compensation to compute motion vectors and block error signals. For each block M_N in an image frame N , a search is performed across the image of a next successive video frame $N+1$ or immediately preceding image frame $N-1$ (i.e., bi-directionally) to identify the most similar respective blocks M_{N+1} or M_{N-1} . The location of the most similar block relative to the block M_N is encoded with a motion vector (DX,DY). The motion vector is then used to compute a block of predicted sample values. These predicted sample values are compared with block M_N to determine the block error signal. The error signal is compressed using a texture coding method such as discrete cosine transform (DCT) encoding.

Object-based video coding techniques have been proposed as an improvement to the conventional frame-based coding standards. In object-based coding, arbitrary shaped image features are separated from the frames in the video sequence using a method called "segmentation." The video objects or "segments" are coded independently. Object-based coding can improve the compression rate because it increases the interframe correlation between video objects in successive frames. It is also advantageous for variety of applications that require access to and tracking of objects in a video sequence.

In the object-based video coding methods proposed for the MPEG-4 standard, the shape, motion and texture of video objects are coded independently. The shape of an object is represented by a binary or alpha mask that defines the boundary of the arbitrary shaped object in a video frame. The motion of an object is similar to the motion data of MPEG-2, except that it applies to an arbitrary-shaped image of the object that has been segmented from a rectangular frame. Motion estimation and compensation is performed

on blocks of a “video object plane” rather than the entire frame. The video object plane is the name for the shaped image of an object in a single frame.

The texture of a video object is the image sample information in a video object plane that falls within the object’s shape. Texture coding of an object’s image samples and error signals is performed using similar texture coding methods as in frame-based coding. For example, a segmented image can be fitted into a bounding rectangle formed of macroblocks. The rectangular image formed by the bounding rectangle can be compressed just like a rectangular frame, except that transparent macroblocks need not be coded. Partially transparent blocks are coded after filling in the portions of the block that fall outside the object’s shape boundary with sample values in a technique called “padding.”

Frame-based coding techniques such as MPEG-2 and H26X and object-based coding techniques proposed for MPEG-4 are similar in that they perform intraframe and interframe coding on macroblocks. Each macroblock includes a series of overhead parameters that provide information about the macroblock. As an example, Fig. 1 shows macroblock parameters used in the header of an interframe macroblock. The COD parameter (10) is a single bit indicating whether the interframe macroblock is coded. In particular, this bit indicates whether or not the encoded macroblock includes motion data and texture coded error data. In cases where the motion and error signal data are zero, the COD bit reduces the information needed to code the macroblock because only a single bit is sent rather than additional bits indicating that the motion vector and texture data are not coded.

In addition to the COD bit, the coding syntax for macroblocks includes coded block parameters (CBP) indicating whether the coded transform coefficients for chrominance and luminance are transmitted for the macroblock. If the transform coefficients are all zero for a block, then there is no need to send texture data for the block. The Coded Block Parameters for chrominance (CBPC) are two bits indicating whether or not coded texture data is transmitted for each of the two chrominance blocks.

The CBPC bits are encoded along with another flag that provides information about the type of quantization for the macroblock. These flags are combined to form a

parameter called MCBPC (12), and MCBPC is entropy coded using an entropy coding method such as Huffman or arithmetic coding.

The parameter called the AC_Pred_flag (14) is a flag indicating whether AC prediction is used in the macroblock.

5 The Coded Block Pattern for luminance (CBPY) (16) is comprised of four bits indicating whether or not coded texture data is transmitted for each of the four luminance blocks. Like the MCBPC parameter, the CBPY flags are also entropy coded using either Huffman or arithmetic coding.

10 After the CBPY parameter, the macroblock includes encoded motion vector data (shown as item 18 in Fig. 1). Following the motion vector data, the "block data" represents the encoded texture data for the macroblock (shown as block data 20 in Fig. 1).

15 One drawback of the coding approach illustrated in Fig. 1 is that it codes CBPC and CBPY flags separately, and therefore, does not exploit the correlation between these parameters to reduce the macroblock overhead. In addition, it does not take advantage of the spatial dependency of the coded block parameters.

SUMMARY

20 The invention provides an improved method of coding the macroblock header parameters in video coding applications. One aspect of the invention is a coding method that exploits the correlation between the coded block parameters by jointly coding all of the coded block parameters with a single variable length code. Another aspect of the invention is a coding method that takes advantage of the spatial dependency between the coded block patterns of neighboring blocks.

25 In an implementation of the invention, the coded block parameters for luminance and chrominance in a macroblock are formed into a single, combined parameter for the macroblock. The combined parameter is assigned a variable length code from a variable length coding table. The coding table is trained based on a target bit rate (e.g., low bit rate Internet applications) and a target class of video content (e.g., talking head video).
30 By jointly coding the luminance and chrominance values, the encoder exploits the correlation between these parameters in the macroblock.

To improve the coding efficiency further, the implementation uses prediction to take advantage of the spatial dependency of the coded block parameters of neighboring blocks. Before assigning the variable length code to the combined parameter, some of the coded block parameters are predicted from neighboring blocks. For intra frame
5 macroblocks, for example, the encoder computes a spatially predicted value for each coded block parameter for luminance. This spatially predicted parameter forms part of the combined parameter for the macroblock.

Additional features and advantages of the invention will become more apparent from the following detailed description and accompany drawings of an implementation of
10 the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram illustrating an example of a macroblock header used in a standard video coding process.

15 Fig. 2 is a block diagram of a video coder.

Fig. 3 is a block diagram of a video decoder.

Fig. 4 is a diagram illustrating an example of an improved macroblock header in which the coded block parameters for chrominance and luminance are jointly coded with a single variable length code.

20 Fig. 5 is a flow diagram illustrating how an implementation of the invention computes a single variable length code for the coded block parameters of I and P frame macroblocks.

Fig. 6 is a diagram illustrating four macroblocks, and their corresponding luminance (Y) blocks.

25 Fig. 7 is a diagram showing an example of the vertical and horizontal gradients of coded block parameter values for selected luminance blocks in Fig. 6.

Fig. 8 is a flow diagram illustrating a method for computing a predictor for coded block parameters.

Fig. 9 is a diagram of a computer system that serves as an operating environment
30 for a software implementation of the invention.

DETAILED DESCRIPTION

Introduction

The first section below provides a description of a video encoder and decoder. Subsequent sections describe how to improve the coding of macroblock header parameters by exploiting the correlation between CBPC and CBPY parameters and taking advantage of the spatial dependency of coded block parameters of neighboring blocks.

Useful in both frame-based and object-based video coding, the invention improves the coding of macroblock parameters, whether the macroblocks are components of arbitrary video objects segmented from a sequence of frames or of rectangular shaped image frames. Object-based coding uses similar motion and texture coding modules as used in frame-based coding. In addition, object-based coders also include shape coding modules. The block syntax relevant to the invention is similar in both frame-based and object-based coding. While the encoder and decoder described in the next section are object-based, they provide a sufficient basis for explaining how to implement the invention in both frame-based and object-based coding schemes.

Description of an Example Encoder and Decoder

Fig. 2 is a block diagram illustrating an implementation of an object-based video encoder. The input 30 to the encoder includes a series of objects, their shape information and bounding rectangles. The shape information, therefore, is available before the encoder codes texture or motion data. Frame-based coding differs in that the entire frame is coded without shape information.

The shape coding module 32 receives the definition of an object including its bounding rectangle and extends the bounding rectangle to integer multiples of macroblocks. The shape information for an object comprises a mask or "alpha plane." The shape coding module 32 reads this mask and compresses it, using for example, a conventional chain coding method to encode the contour of the object.

Motion estimation module 34 reads an object including its bounding rectangle and a previously reconstructed image 36 and computes motion estimation data used to predict the motion of the object from one frame to another. After identifying the macroblocks in the current object image, the motion estimation module 34 searches for the most similar

macroblock in the reconstructed image for each macroblock in the current object image to compute the motion data for each macroblock. The specific format of the motion data from the motion estimation module 34 can vary depending on the motion estimation method used. The implementation described below computes a motion vector for each
5 macroblock, which is consistent with current MPEG and H26X formats.

The motion compensation module 38 reads the motion vectors computed by the motion estimation module and the previously reconstructed image 36 and computes a predicted image for the current frame. The encoder finds the difference between the image sample values in the input image block as specified in the input 30 and the
10 corresponding sample values in the predicted image block as computed in the motion compensation module 38 to determine the error signal for the macroblock.

Texture coding module 40 compresses this error signal for inter-frame coded objects and compresses image sample values for the object from the input data stream 30 for intra-frame coded objects. The feedback path 42 from the texture coding module 40
15 represents the decoded error signal. The encoder uses the error signal macroblocks along with the predicted image macroblocks from the motion compensation module to compute the previously reconstructed image 36.

The texture coding module 40 codes blocks of intra-frame and error signal data for an object using any of a variety of still image compression techniques. Example
20 compression techniques include transform-based techniques such as DCT and wavelet coding as well as other conventional image compression methods such as LaPlacian Pyramid coding.

The bitstream of the compressed video sequence includes the shape, motion and texture coded information from the shape coding, motion estimation, and texture coding
25 modules. Multiplexer 44 combines and formats this data into the proper syntax and outputs it to the buffer 46.

While the encoder can be implemented in hardware or software, it is most likely implemented in software. In a software implementation, the modules in the encoder represent software instructions stored in memory of a computer and executed in the
30 processor and the video data stored in memory. A software encoder can be stored and distributed on a variety of conventional computer readable media. In hardware

implementations, the encoder modules are implemented in digital logic, preferably in an integrated circuit. Some of the encoder functions can be optimized in special-purpose digital logic devices in a computer peripheral to off-load the processing burden from a host computer.

5 Fig. 3 is a block diagram illustrating a decoder for an object-based video coding method. A demultiplexer 60 receives a bitstream 62 representing a compressed video sequence and separates shapes, motion and texture encoded data on an object by object basis. Shape decoding module 64 decodes the shape or contour for the current object being processed. To accomplish this, it employs a shape decoder that implements the
10 inverse of the shape encoding method used in the encoder of Fig. 2. The resulting shape data is a mask, such as a binary alpha plane or gray scale alpha plane representing the shape of the object.

 The motion decoding module 66 decodes the motion information in the bitstream. The decoded motion information includes motion data such as motion vectors for
15 macroblocks blocks or geometric transform coefficients, depending on the type of estimation method used in the encoder. The motion decoding module 66 provides this motion information to the motion compensation module 68, and the motion compensation module 68 applies the motion data to previously reconstructed object data 70.

 The texture decoding module 74 decodes error signals for inter-frame coded
20 texture data and an array of color values for intra-frame texture data and passes this information to a module 72 for computing and accumulating the reconstructed image. For inter-frame coded objects, this module 72 applies the error signal data to the predicted image output from the motion compensation module to compute the reconstructed object for the current frame. For intra-frame coded objects the texture
25 decoding module 74 decodes the image sample values for the object and places the reconstructed object in the reconstructed object module 72. Previously reconstructed objects are temporarily stored in object memory 70 and are used to construct the object for other frames.

 Like the encoder, the decoder can be implemented in hardware, software or a
30 combination of both. In software implementations, the modules in the decoder are software instructions stored in memory of a computer and executed by the processor and

video data stored in memory. A software decoder can be stored and distributed on a variety of conventional computer readable media. In hardware implementations, the decoder modules are implemented in digital logic, preferably in an integrated circuit. Some of the decoder functions can be optimized in special-purpose digital logic devices
5 in a computer peripheral to off-load the processing burden from a host computer.

Improved Coding of Macroblock Overhead

The invention includes innovations that improve the coding of macroblock header parameters. One innovation is a method for coding the coded block parameters to exploit
10 the correlation between CBPC and CBPY. This innovation is implemented by jointly coding a combined CBPC and CBPY parameter with a single variable length code. Another innovation further improves coding efficiency of the header parameters by exploiting the spatial dependency of the coded block parameters. In particular, coded block parameters are more efficiently compressed by predicting them from the parameter
15 of neighboring blocks.

Figure 4 is a diagram illustrating the header block parameters computed by an implementation of the invention. Like the header information shown in Fig.1, this header block includes a COD parameter 80, an AC_Pred_flag 82, motion vector data (MV 84) and block data 86. Unlike the header in Fig. 1, MCBPC and CBPY parameters are
20 jointly coded with a single variable length code, called MBCBPCY 88. This code combines coded block parameters for chrominance and luminance, as well as the flag for macroblock type.

Figure 5 is a flow diagram illustrating how the implementation generates a variable length code for Intra (I) frames and predicted (P) frames. In this particular
25 implementation, the header blocks for I and P frames are coded differently. For I frames, the encoder performs the additional step of predicting the coded block parameters for luminance before selecting the variable length code. It is also possible to use prediction for P frames. However, prediction does not improve coding efficiency significantly in P frames, and in some cases, can even decrease coding efficiency.

30 The goal of using prediction for coded block parameters is to produce as many zero values for these parameters as possible. By making the values mostly zero, the

encoder reduces the variance of the coded block parameters. The process of training the variable length coding table can then favor the zero value, which improves coding efficiency. In P frames, especially in low bit rate applications, the coded block parameters are mostly zero before prediction. As such, prediction does not tend to
5 increase the number of zero values, and sometimes, it even decreases the number of zero values. Therefore, the implementation shown in Fig. 5 does not use prediction for P frames.

For P frames, the encoder begins by finding the coded block parameters for luminance and chrominance as shown in step 100. These block parameters are each a
10 single bit indicating whether a corresponding block is texture coded. The coded block parameters are computed in the texture coding module (40 in Fig. 2), which sets a coded block flag for each block that has non-zero encoded texture values. Conversely, the value of the coded block parameter for a block in which the texture values are all zero (or so close to zero as to be negligible) is zero.

15 Since there are two blocks for chrominance (one each for the 8 by 8 pixel U and V blocks) and four blocks for luminance (one each for the four 8 by 8 blocks) in the macroblock, the combined parameter for the coded block pattern is a total of six bits. Combining this 6 bit number with the single bit for macroblock type, the encoder forms a 7 bit number as shown in step 102. The macroblock type indicates whether the
20 macroblock is for an I or P frame.

Once the combined MBCBPCY is formed, the combined parameter is looked up in a variable length coding table to find a corresponding variable length code associated with the parameter as shown in step 104. The encoder assigns a single variable length code to the combined parameter MBCPCY.

25 The coding table in the implementation table is a Huffman coding table. The table is preferably trained based on the target rate and target scenario. Table 1 below is a Variable Length Coding (VLC) table obtained for a low bit rate "talking head" scenario. For each macroblock in a P frame, the combined MBCBPCY information is coded using the code word for the corresponding entry in this table.

Table 1: VLC Table for Coded block pattern of chrominance and luminance for P picture

Index	MB type	CBPCY Y(1234)UV	Number of bits	Code
0	I	000000	7	1000000
1	I	000001	13	1001111001001
2	I	000010	12	100111111101
3	I	000011	15	000000111111100
4	I	000100	12	100111111100
5	I	000101	18	000000101010000011
6	I	000110	17	10010110100110100
7	I	000111	16	1000001110111100
8	I	001000	12	100000111010
9	I	001001	17	00000011111111000
10	I	001010	16	0000001111111101
11	I	001011	16	0000001111111111
12	I	001100	13	0000001111001
13	I	001101	18	000000101010000010
14	I	001110	16	1001011010011101
15	I	001111	16	0000001010100100
16	I	010000	12	1001011111000
17	I	010001	17	00000010101000011
18	I	010010	15	100000111011111
19	I	010011	17	00000011111111001
20	I	010100	13	1001011110011
21	I	010101	18	100101101001101011
22	I	010110	18	100101111011111001
23	I	010111	16	0000001111111010
24	I	011000	14	10000011101110
25	I	011001	20	1001011010011010101 1
26	I	011010	16	1001011010011100
27	I	011011	18	100101111011111000
28	I	011100	13	1001011010010
29	I	011101	18	000000101010000101
30	I	011110	16	1001011010011110
31	I	011111	15	100101111001000
32	I	100000	12	000000111101
33	I	100001	17	10010111101111111
34	I	100010	16	0000001010100010
35	I	100011	16	1001011010011111
36	I	100100	14	10010111101110
37	I	100101	21	1001011010011010101 01
38	I	100110	17	10010111101111101
39	I	100111	17	10010111101111110
40	I	101000	12	100111100101
41	I	101001	18	000000101010000001
42	I	101010	19	1001011010011010100
43	I	101011	16	1000001110111101
44	I	101100	13	0000001111000
45	I	101101	16	1001011010011011

Index	MB type	CBPCY Y(1234)UV	Number of bits	Code
46	I	101110	16	0000001111111110
47	I	101111	16	0000001010100101
48	I	110000	13	0000001111110
49	I	110001	18	000000101010000000
50	I	110010	16	0000001010100011
51	I	110011	16	0000001111111011
52	I	110100	13	1000001110110
53	I	110101	18	000000101010000100
54	I	110110	15	000000101010011
55	I	110111	15	100101111001001
56	I	111000	13	0000001010101
57	I	111001	21	1001011010011010101 00
58	I	111010	15	100101111011110
59	I	111011	14	10010111100101
60	I	111100	10	1001011011
61	I	111101	15	100101101001100
62	I	111110	12	100101101011
63	I	111111	12	100101101010
64	P	000000	2	01
65	P	000001	7	0000000
66	P	000010	6	100110
67	P	000011	9	100101011
68	P	000100	3	111
69	P	000101	10	1000001111
70	P	000110	9	000000100
71	P	000111	12	000000101000
72	P	001000	3	110
73	P	001001	10	1000001010
74	P	001010	9	100101000
75	P	001011	12	000000101011
76	P	001100	5	10001
77	P	001101	11	00000011011
78	P	001110	9	100111010
79	P	001111	11	10011111111
80	P	010000	4	0011
81	P	010001	10	1001110111
82	P	010010	9	100000110
83	P	010011	12	100000111001
84	P	010100	4	1011
85	P	010101	10	1001111011
86	P	010110	9	100101100
87	P	010111	11	10010111111
88	P	011000	6	001001
89	P	011001	12	000000110101
90	P	011010	10	1001111110
91	P	011011	13	1001111001000
92	P	011100	6	000001
93	P	011101	11	10010101010

Index	MB type	CBPCY Y(1234)UV	Number of bits	Code
94	P	011110	10	1000001000
95	P	011111	12	000000101001
96	P	100000	4	0001
97	P	100001	10	1001010100
98	P	100010	9	100101110
99	P	100011	12	100000111000
100	P	100100	6	100100
101	P	100101	11	10011110011
102	P	100110	10	1001110110
103	P	100111	13	1001011110110
104	P	101000	5	00001
105	P	101001	10	1001111010
106	P	101010	9	100111110
107	P	101011	12	000000111110
108	P	101100	6	001000
109	P	101101	11	10000010011
110	P	101110	10	0000001100
111	P	101111	11	10010111110
112	P	110000	5	10100
113	P	110001	11	10000010010
114	P	110010	10	1001010011
115	P	110011	12	100101111010
116	P	110100	6	100001
117	P	110101	11	10010101011
118	P	110110	10	1000001011
119	P	110111	12	000000110100
120	P	111000	5	10101
121	P	111001	10	1001111000
122	P	111010	10	1001010010
123	P	111011	12	100101101000
124	P	111100	5	00101
125	P	111101	10	0000001011
126	P	111110	8	10011100
127	P	111111	10	0000001110

In the implementation shown in Fig. 5, I frames are coded differently than P frames in that the encoder uses prediction to exploit the spatial dependency of the coded block parameters. For each macroblock, the encoder begins by getting the coded block parameters for chrominance and luminance as shown in step 106.

Next, the encoder computes the predictor of the coded block parameters for luminance. In this particular implementation, the encoder only uses prediction for the CBPY parameters. However, the same prediction method could also be used to predict the coded block parameters for chrominance. In the case of chrominance, the prediction

is computed based on 8 by 8 pixel chrominance blocks in neighboring macroblocks rather than the neighboring 8 by 8 pixel luminance blocks, which may be in the same macroblock or a neighboring macroblock. Since each macroblock has four luminance blocks, the neighboring blocks for a given luminance block may come from the same or neighboring macroblock. For prediction involving chrominance blocks, the neighboring blocks come from neighboring macroblocks.

The encoder performs spatial prediction on coded block parameters. First, it looks at the coded block parameters for neighboring blocks to determine whether the value of the block parameter is likely to change from a neighboring block to the current block of interest. If the location of a block representing the smallest change in the coded block parameter can be identified (i.e. the lowest spatial gradient in the coded block parameters), then the coded block parameter for the block at this location is used as the predictor. Otherwise, it does not matter which neighbor is chosen as the predictor and one is merely selected. A specific example of selecting the predictor is described and illustrated in more detail with reference to Figs. 6-8 below.

In the next step 110, the encoder computes a predicted value for the coded block parameters. The predicted value represents the change in the coded block parameter for the predictor block and the current block. To compute the predicted value, the encoder performs a bitwise exclusive OR (XOR) on the predicted value and current block value. The resulting vector, called CBPCY_XOR is then assigned a variable length code from a Huffman table. The encoder looks up the entry for CPCY_XOR in the table and finds the corresponding variable length code. Table 2 below shows the VLC table used to code predicted CBPCY values for I frames in the implementation.

Table 2 : VLC Table for Coded block pattern of chrominance and luminance for I picture

Index	CBPCY_XOR Y(1234)UV	Number of bits	Code
0	000000	1	1
1	000001	6	010111
2	000010	5	01001
3	000011	5	00101
4	000100	5	00110
5	000101	9	001000111
6	000110	7	0100000
7	000111	7	0010000
8	001000	5	00010

Index	CBPCY_XOR Y(1234)UV	Number of bits	Code
9	001001	9	001111100
10	001010	7	0111010
11	001011	7	0011101
12	001100	6	000010
13	001101	9	011101100
14	001110	8	01110111
15	001111	8	00000000
16	010000	5	00011
17	010001	9	010110111
18	010010	7	0101100
19	010011	7	0010011
20	010100	6	000001
21	010101	10	0101101000
22	010110	8	01000110
23	010111	8	00111111
24	011000	6	011110
25	011001	13	0011100010010
26	011010	9	010110101
27	011011	8	01000010
28	011100	7	0100010
29	011101	11	00111000101
30	011110	10	0100011110
31	011111	9	010000111
32	100000	4	0110
33	100001	9	000000011
34	100010	7	0011110
35	100011	6	011100
36	100100	7	0010010
37	100101	12	001110001000
38	100110	9	001000100
39	100111	9	001110000
40	101000	6	011111
41	101001	11	01000111110
42	101010	8	00111001
43	101011	9	010001110
44	101100	7	0000001
45	101101	11	00111000110
46	101110	9	010110110
47	101111	9	001000101
48	110000	6	010100
49	110001	11	01000111111
50	110010	9	001111101
51	110011	9	000011000
52	110100	7	0000111
53	110101	11	00111000111
54	110110	9	010000110
55	110111	9	000011001
56	111000	6	010101
57	111001	10	0111011011
58	111010	9	000000010

Index	CBPCY_XOR Y(1234)UV	Number of bits	Code
59	111011	9	001000110
60	111100	8	00001101
61	111101	13	0011100010011
62	111110	10	0111011010
63	111111	10	0101101001

Figs. 6-8 illustrate the spatial prediction performed in the encoder in more detail.

Fig. 6 is a diagram showing four neighboring macroblocks (top left-120, top right-122, lower left-124, and lower right-126). The following example focuses on the lower right block, which is circled. Each of the macroblocks includes four 8 by 8 pixel blocks for luminance labeled as Y1, Y2, Y3 and Y4.

As an example, consider the top left luminance block Y1 for macroblock 126. The blocks used to compute the predictor are surrounded by a dashed line 128. The block of interest is Y1 (labeled as block 130a), and the blocks used to compute the predictor are the neighboring blocks labeled as 132a, 134a, and 136a.

To give a specific example, Fig. 7 shows values of the coded block pattern parameters for each of the blocks within the dashed line of Fig. 6. The reference numbers 130b, 132b, 134b and 136b correspond to the blocks 130a, 132a, 134a and 136a of Fig. 6, respectively. The spatial gradients of the neighboring coded blocks parameters are used to select the predictor. In particular, the vertical gradient is computed from the coded block parameters of the top-left and left neighboring blocks (136a, 132a, shown circled 140 in Fig. 7). The horizontal gradient is computed from the coded block parameters of the top-left and top neighboring blocks (136a, 130a, shown circled 142 in Fig. 7).

Fig. 8 is a flow diagram illustrating the steps for finding the predictor. First, the encoder finds the vertical and horizontal gradients. Each is computed as the exclusive OR of the coded block parameters shown circled in Fig. 7 (140 is the vertical gradient and 142 is the horizontal gradient). Next, the encoder compares the gradient values. If the gradients are not the same, the encoder selects the predictor as the value assigned to the block in the direction of the lower gradient. In the example shown in Fig. 7, the vertical gradient is zero, while the horizontal gradient is one. Thus, the direction of the lower gradient is up. As such, the value of the coded block parameter for block 134a is

used as the predictor because it is located in the “up” direction relative to the block of interest.

Whether or not prediction is used to modify the coded block parameters, the end result is a single variable length code representing all of the coded block parameters for the macroblock. Since I and P frames are coded differently in the implementation, the decoder treats the macroblocks for these frames differently. For P frames, the decoder uses VLC table 1 to look up the single variable length code and find the corresponding entry that stores the combined parameter representing the coded block parameters for luminance and chrominance. For I frames, the decoder uses VLC table 2 to look up the single variable length code and find the corresponding entry that stores the combined parameter representing coded block parameters for luminance and chrominance. In both I and P frames, the texture decoding module (block 74 in Fig. 3) uses the coded block parameters to determine whether the texture data for the corresponding block needs to be decoded. The decoder skips texture decoding for blocks having a coded block parameter of zero.

In cases where the coded block parameters are also predicted, the decoder uses the previously decoded block parameters from the neighboring blocks to compute the coded block parameter for the current block of interest. First, the decoder computes the location of the predictor block based on the spatial gradients in the same manner as in the encoder. Next, it computes the value of the coded block parameter for the current block by computing the exclusive OR of the decoded value and the coded block parameter of the predictor block (the exclusive OR operator has the following property: $X \text{ XOR } Y = Z$; $Z \text{ XOR } X = Y$). After this inverse prediction stage, the texture decoder then uses the coded block parameter to determine whether to skip decoding the texture for the block.

Brief Overview of a Computer System

Figure 9 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the invention may be implemented. Although the invention or aspects of it may be implemented in a hardware device, the encoder and decoder described above are implemented in computer-executable instructions organized in program modules. The program modules include the

routines, programs, objects, components, and data structures that perform the tasks and implement the data types described above.

While Fig. 9 shows a typical configuration of a desktop computer, the invention may be implemented in other computer system configurations, including hand-held
5 devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, and the like. The invention may also be used in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote
10 memory storage devices.

Figure 9 illustrates an example of a computer system that serves as an operating environment for the invention. The computer system includes a personal computer 920, including a processing unit 921, a system memory 922, and a system bus 923 that interconnects various system components including the system memory to the processing
15 unit 921. The system bus may comprise any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using a bus architecture such as PCI, VESA, Microchannel (MCA), ISA and EISA, to name a few. The system memory includes read only memory (ROM) 924 and random access memory (RAM) 925. A basic input/output system 926 (BIOS), containing the basic routines that
20 help to transfer information between elements within the personal computer 920, such as during start-up, is stored in ROM 924. The personal computer 920 further includes a hard disk drive 927, a magnetic disk drive 928, e.g., to read from or write to a removable disk 929, and an optical disk drive 930, e.g., for reading a CD-ROM disk 931 or to read from or write to other optical media. The hard disk drive 927, magnetic disk drive 928,
25 and optical disk drive 930 are connected to the system bus 923 by a hard disk drive interface 932, a magnetic disk drive interface 933, and an optical drive interface 934, respectively. The drives and their associated computer-readable media provide nonvolatile storage of data, data structures, computer-executable instructions (program code such as dynamic link libraries, and executable files), etc. for the personal computer
30 920. Although the description of computer-readable media above refers to a hard disk, a removable magnetic disk and a CD, it can also include other types of media that are

readable by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, and the like.

A number of program modules may be stored in the drives and RAM 925, including an operating system 935, one or more application programs 936, other program modules 937, and program data 938. A user may enter commands and information into the personal computer 920 through a keyboard 940 and pointing device, such as a mouse 942. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 921 through a serial port interface 946 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port or a universal serial bus (USB). A monitor 947 or other type of display device is also connected to the system bus 923 via an interface, such as a display controller or video adapter 948. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

The personal computer 920 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 949. The remote computer 949 may be a server, a router, a peer device or other common network node, and typically includes many or all of the elements described relative to the personal computer 920, although only a memory storage device 950 has been illustrated in Figure 9. The logical connections depicted in Figure 9 include a local area network (LAN) 951 and a wide area network (WAN) 952. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the personal computer 920 is connected to the local network 951 through a network interface or adapter 953. When used in a WAN networking environment, the personal computer 920 typically includes a modem 954 or other means for establishing communications over the wide area network 952, such as the Internet. The modem 954, which may be internal or external, is connected to the system bus 923 via the serial port interface 946. In a networked environment, program modules depicted relative to the personal computer 920, or portions thereof, may be stored in the remote memory storage device. The network

connections shown are merely examples and other means of establishing a communications link between the computers may be used.

Conclusion

5 While the invention has been illustrated using specific implementation as an example, the scope of the invention is not limited to the specific implementation. For example, it is possible to use spatial prediction for both chrominance and luminance blocks using similar techniques. In addition, spatial prediction may be used for coding the coded block parameters for both intra and predicted frames. The implementation uses
10 Huffman tables to generate variable length codes. In fact, a variety of entropy coding methods may be used to generate a variable length code for each combined coded block parameter. For instance, various forms of arithmetic and/or run length encoding may be used. Each of these coding methods assign longer codes to input signals that occur less frequently while assigning shorter coders to more frequent input signals. As noted above,
15 the coding methods for improving the efficiency of macroblock headers can be applied to frame-based and object based coding methods.

 In view of the many possible implementations of the invention, it should be recognized that the implementation described above is only an example of the invention and should not be taken as a limitation on the scope of the invention. Rather, the scope of
20 the invention is defined by the following claims. We therefore claim as our invention all that comes within the scope and spirit of these claims.